Off-line Computing Capabilities at LBNL/NERSC

Doug Olson, LBNL

STAR Collaboration Meeting

2 August 1999, BNL

Outline

- Background of PDSF for STAR (http://pdsf.nersc.gov/)
- NERSC
- PDSF history
- Present configuration of STAR cluster
- Primary activities at PDSF/NERSC

RHIC Off-line Computing Model

(System Will Be Partitioned By Experiment)

Conceived late 1995

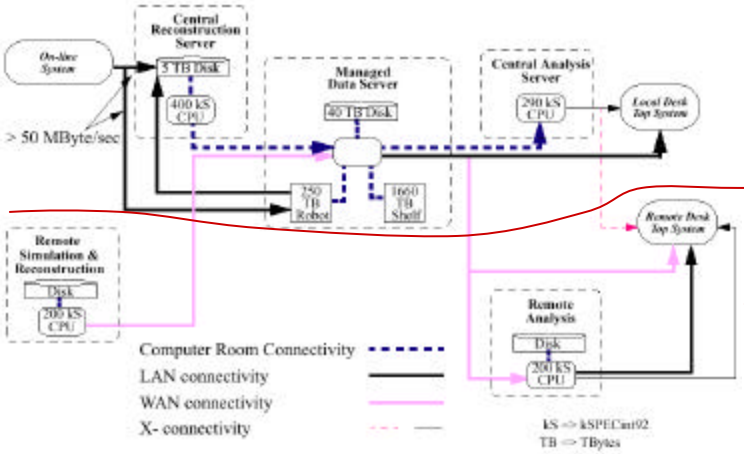

BNL onsite
- main data center
- 50% CPU

Remote
- modest data
- 50% CPU
- leverage additional facilities

Offline Computing at RHIC, Kumar et al. (ROCC), Feb. 1996

Implementation of RHIC Computing Model

Extremely successful 4 years since conception!

Lots of participation and interest beyond BNL!

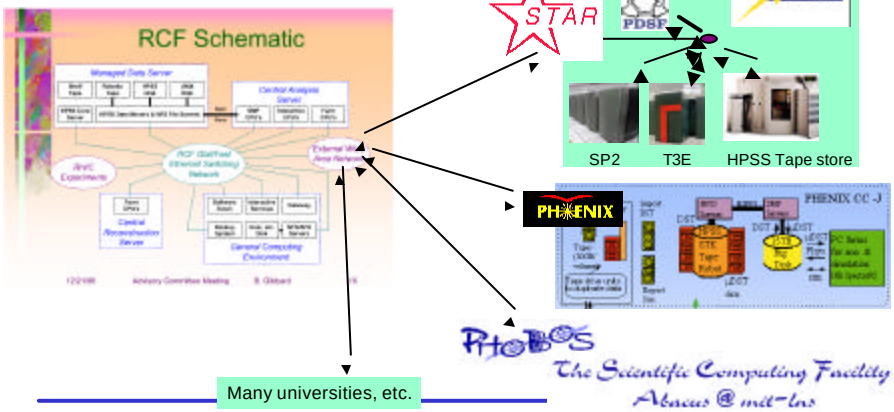

Many universities, etc.

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Advancing Computational Science of Scale—Producing Real Results

- One of the nations most powerful unclassified computing resources.
  - NERSC is funded by the Department of Energy, Office of Science, and is part of Computing Sciences at Lawrence Berkeley National Laboratory.
- Production facilities
  - 1 Teraflop (late 1999), 3 Teraflops (late 2000) (Cray & IBM)
  - 600 Terabytes robotic tape storage (HPSS)
  - Parallel Distributed Systems Facility (PDSF) linux farm (more later…)
- Large R&D program
- Adjacent to ESnet headquarters

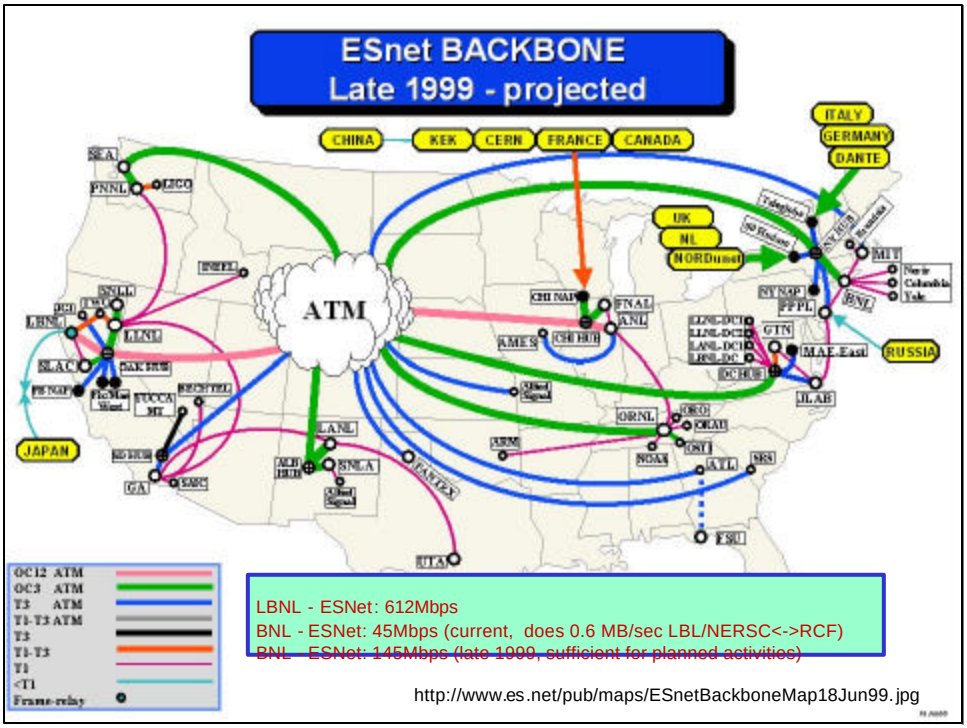

NERSC PDSF - Mar. 1999

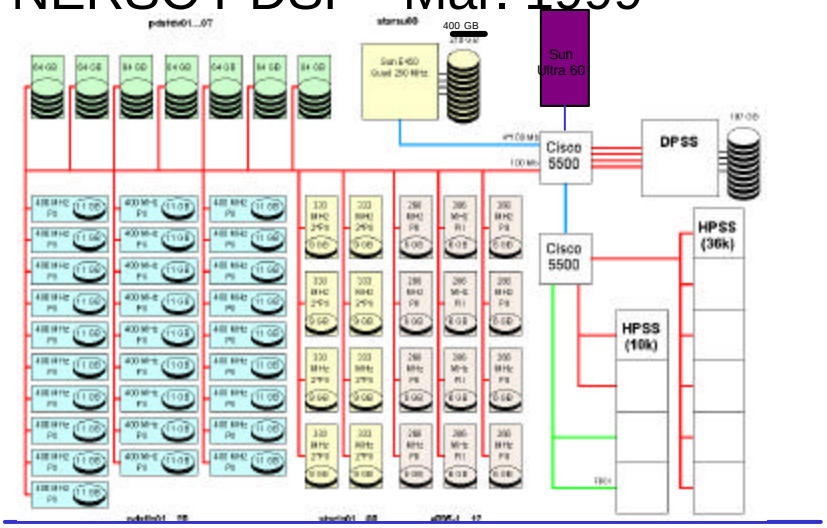

Overview of PDSF for STAR

- Developed in response to RHIC computing plan
- Simulations and analysis capability for STAR
  - at LBNL under control of STAR collaboration
  - processor farm and disk space coupled to NERSC mass storage
  - ESnet interface to RCF
- Goal of PDSF for STAR
  - Provide scaleable system for simulations and downstream analysis for the STAR collaboration
  - Long term: ~= STAR share of RCF cpu
    ~ 6K SPECint95, 40 TB/year robotic tape, 15 TB disk
  - Buildup now funded by LBNL Nuclear Science Division and NERSC
  - Expect DOE funding beginning in FY00

via NERSC allocation

History of PDSF hardware

- Arrived from SSC (RIP), May 1997
- October '97
  - 32 SUN Sparc 10, 32 HP 735
  - 2 SGI data vaults
- January 1998
  - added 12 single-cpu Intel (LBNL/E895)
- June 1998
  - added SUN E450, 8 dual-cpu Intel (LBNL/STAR)
- October 1998
  - added 16 single-cpu Intel (NERSC)
  - added 500 GB network disk (NERSC)
  - subtracted SUN, HP
  - subtracted 160 GB SGI data vaults
- January 1999
  - added 12 single-cpu Intel (NERSC)
  - added 200 GB disk (LBNL/STAR)
- July 1999
  - added 900 GB disk (LBNL/STAR)
- September 1999
  - adding 18 dual-cpu Intel (LBNL/STAR)
  - adding 1 TB network RAID disk (LBNL/STAR)
- **Totals (Sept. 1999)**
  - **1.2 TB disk on SUN E450**
  - **1.5 TB network disk**
  - **1500 SPECint95**
  - **archival storage (via NERSC allocation of 600TB system)**

Comparison of RCF to PDSF

(Oct. 1999)

	RCF	PDSF	PDSF/RCF (%)
Total CPU (SPECint95)	6800	1500	22%
STAR share	**2473**	**1200**	**49%**
Total Disk (TB)	8.1	2.7	33%
STAR share	**2.95**	**2.5**	**85%**

PDSF Software

- STAR off-line environment
  - Linux, Solaris/Sparc
- Packages
  - AFS, CERNlib, DPSS, EGCS, Framereader, GNU, HPSS, itcl, Java, kai, LSF, Objectivity, Orbacus, Orbix, PSDF Login, Perl, pgi, pvm, Root, SSH, SUN Workshop, TCL, TK
  - >250 software modules for Solaris/Sparc

Current PDSF staff

- 2.5 FTE NERSC + 0.5 FTE NSD
  - Craig Tull - group leader
  - Thomas Davis (1) - PDSF project lead
  - Cary Whitney (1), John Milford (1/2)
  - Dieter Best (1/4 NSD), Qun Li (1/4 NSD)

Supercomputer time for Gstar production

- FY99
  - Used 250K hours of Cray T3E at PSC + NERSC
  - Amortized over 4 months: ~ 1000 SPECint95
- FY2000
  - Have requested 300K hours at NERSC on T3E and new IBM SP2 system.

STAR Activities at PDSF/NERSC

- Interactive s/w development on PDSF
- Individual data analysis (ROOT, Gstar, …)
  - now: 58 STAR users, 27 from LBL
- Gstar production (T3E, SP2 systems)
- Detector response simulation
- BFC reconstruction of simulated data
  - full simulation, and embedding studies
- Mirror subset of real data for analysis

Planned for FY2000

Summary

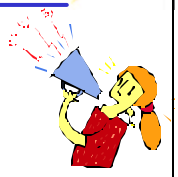

- Simulations/Analysis capability at LBNL for STAR is growing
  - Oct. 99: 49% cpu, 85% disk of STAR @ RCF
- Currently 58 STAR users, 31 from outside LBNL
- PDSF receiving strong support
  - ($, people) from LBNL Nuclear Science Division and NERSC
- In FY00, focus on:
  - Simulation - detector response, reconstruction, embedding
  - Analysis of real data
- Learn more about PDSF & become a user! http://pdsf.nersc.gov